

High-Resolution AMR Simulations of Confined Explosions*

John Bell

William Crutchfield

Charles Rendleman

Vincent Beckner

Mike Lijewski

Lawrence Berkeley National Laboratory
Berkeley, CA 94720

Allen Kuhl

Lawrence Livermore National Laboratory
Livermore, CA 94550

June 2, 2000

Abstract

We are developing end-to-end simulation capabilities for modeling explosion effects related to counter-proliferation and counter-terrorism scenarios in support of the Defense Threat Reduction Agency (DTRA). These problems include explosions in buried chamber systems for military applications and explosions in or near buildings for civilian applications. Our approach, which is based on high-resolution adaptive numerical methods (developed under DTRA's Advanced Computational Methods Program), is designed to provide a capability of performing detailed first-principles calculations of such turbulent explosion fields. The results of such analysis provide vital input to military strategy for Counter-Proliferation measures, aid forensic investigations of terrorist bombing incidents, and suggest methods of protecting US troops abroad and US civilians at home.

In this paper we describe the simulation of a scenario in which an explosion occurs within a chamber containing a canister of combustible material. The explosion results in a rupture in the container followed by the possible detonation/combustion of the contents of the container. The simulation determines the conditions under which the canister's contents are vented to the enclosing environment through an opening at the top of the chamber. Because of the intensive

*Support for this work was provided by the Applied Mathematical Sciences Program of the DOE Office of Mathematics, Information, and Computational Sciences and by the Defense Threat Reduction Agency IACRO98-3017, through the U. S. Department of Energy under Contract No. DE-AC03-76SF00098.

computing requirements for such simulations, we used an Adaptive Mesh Refinement (AMR) method for compressible gas-dynamics, targeted for distributed memory parallel computers. The calculations were performed using the Cray T3E at ARSC and the IBM SP systems at ERDC. The results demonstrate the value of the combination of AMR simulation with advanced parallel computing systems.

1 Introduction

Countering the proliferation of weapons of mass destruction (i.e., chemical, biological or nuclear weapons) is one of the critical mission areas for the Defense Threat Reduction Agency (DTRA). Assessment of explosion effects related to military scenarios provides important inputs to military strategy. For example, if an underground bunker filled with chemical/biological weapons is attacked, DTRA must predict the consequences, such as:

- what is the evolution in space and time of the energy released during the explosion;
- will this energy release destroy the structure;
- will the chemical/biological weapons be breached;
- will the agent be destroyed by the high-temperature environment of the explosion;
- how much of the agent will be ejected from the bunker;
- how will the agent be dispersed in the atmosphere and at what risk to civilian populations?

Countering terrorist bombing incidents (e.g., the recent attack of US troops in the Khobar Towers in Saudi Arabia, and attacks of US civilian populations at the World Trade Center and Oklahoma City) is a developing mission area for DTRA. Application of explosion science to the analysis of terrorist bombing incidents provides valuable technical inputs (e.g., how much explosive is consistent with the observed damage, etc.) to forensic investigations by the DOD and FBI, and suggests methods of protecting US troops abroad and civilians at home.

The goal of this project is to develop the computational tools required for detailed assessment of such explosions. By utilizing advanced computer architectures available in the DOD High Performance Computing Modernization Program combined with sophisticated adaptive numerical methodology, it is now possible to perform detailed, end-to-end numerical simulations of the three-dimensional time-dependent explosion field. A complete simulation capability requires modeling fluid-dynamic phenomena in three different regimes:

Confined Explosion The physical phenomenon is fluid dynamics driven by chemical energy release in geometries such as buried chamber systems or urban areas around buildings and other structures. For the problems of interest to DTRA, this will require

- the representation of both premixed and non-premixed combustion to capture the effects of after-burning;
- strong compressibility effects, due both to shock waves and exothermic expansion; and

- turbulent flow in complex three-dimensional geometries.

An additional numerical issue is the importance of computing on time scales that are long relative to the acoustic time scales (on the order of seconds) to capture distributed exothermicity effects such as after-burning.

Plume Ejection Once the dominant energetic processes are complete and motion on acoustic time scales has equilibrated, the flow is characterized by the dynamics of hot explosion-product gases including hazardous contaminants. The flow consists of a plume of explosion gases emanating from the underground chamber system, and may include complex geometry representing surface structures and local terrain. The primary physical processes that must be represented in numerical calculations include buoyancy and turbulent mixing. In addition, compressibility effects due to choked flow are important (since it controls the flow rate escaping from the pressurized chamber). The goal in this area is to characterize the source to the point where it could be treated as a contaminant passively transported by the fluid flow.

Atmospheric Transport Once turbulent mixing and cooling have essentially equilibrated the explosion products and contaminants with the ambient atmosphere, the long-term fate of contaminants can be modeled as passive transport of the contaminant cloud by the local atmospheric flow. In this regime, we must model the cloud transport over long distances where large-scale effects such as those due to topography, atmospheric stratification, and local meteorology, are dominant.

2 Technical Approach

2.1 Basic Discretization Methodology

Our discretization methods are based on the use of a predictor-corrector temporal discretization [15, 5] for solving advection-diffusion equations with source terms. In this approach, the advection terms are computed using a high-resolution upwind method as in [18, 11, 10, 6], while the diffusion terms are treating semi-implicitly using a Crank-Nicolson type of discretization. This is done in such a way that the overall method is second-order accurate in space and time with a time step limited only by the CFL condition for the advection terms. The linear systems associated with the methodology correspond to standard discretizations of self-adjoint second-order elliptic PDEs for which efficient iterative methods are available. Extension to incompressible and low Mach number flows is performed using a projection formulation [9, 5]. In this approach, the constraint on the divergence of the velocity field is enforced using a Hodge projection thereby transforming the constrained problem into a pure initial value problem, to which a version of our predictor-corrector formalism can be applied.

A fundamental design choice for our methodology is that of the spatial grid structure. We have based our approach on the use of smooth, logically rectangular grids, in which a grid in physical space is the image of a subset of a rectangular lattice in some abstract coordinate space under the action of a smooth mapping. The advantages of the rectangular grid approach are accuracy and efficiency. For example, it is easier to construct discretizations of a given accuracy on a smoothly-varying rectangular grid, due to a cancellation of errors in difference approximations

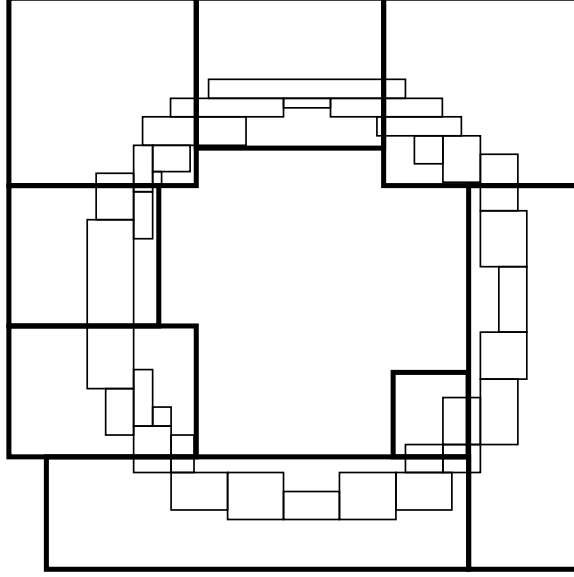


Figure 1: Two levels of refined grids. Grids are properly nested, but may have more than one parent grid. The thick lines represent grids at the coarse level; the thin lines, grids at the fine level.

on smooth meshes. In addition, the resulting calculations are extremely regular, with conventional array data access patterns, thus minimizing the cost of finding the data required to update the solution. In order to accommodate adaptivity and complex boundary geometries, we introduce irregular calculations (i.e. ones that are not performed on rectangular patches) but only on a set of co-dimension one in the problem domain. This makes the cost of the irregular calculations a small fraction of the overall computational cost. In addition, it has a relatively small impact on the overall accuracy of the computation.

2.2 Adaptive Methods

Our adaptive methods are based on the block-structured adaptive mesh refinement (AMR) algorithms of Berger and Oliger [8, 7, 4]. In this approach, the regions to be refined are organized into rectangular patches with several hundred to several thousand grid points per patch (Figure 1). One is thus able to use the rectangular grid methods described above to advance the solution in time; furthermore, the overhead in managing the irregular data structures is amortized over relatively large amounts of floating-point work. Refinement is performed in time as well as in space with each level of refinement using its own time step (subject to the constraint that the time step at a coarse level be an integral multiple of the time step at the next finer level). This leads to an algorithm that is both more efficient and more accurate than one based on using the same time step at all level of refinement.

Initially, AMR was developed for hyperbolic conservation laws; however, the methodology has been extended to a broad range of applications. In particular, we have extended the AMR algorithm to the case of incompressible flows in two and three space dimensions [1, 2, 3] using the second-order projection method. For the advective transport terms in the equations, the algorithm reduces to a variation of the original hyperbolic algorithm. However, the elliptic nature

of the incompressibility constraint introduces additional issues, due to the more complicated nature of the coarse-fine matching conditions for elliptic problems. These issues include identifying the appropriate discretizations of the combined Dirichlet and Neumann matching conditions and constructing an appropriate generalization of those conditions for refinement in time. In addition, the hyperbolic synchronization of momentum transport at coarse-fine boundaries must be decomposed into divergence-free and gradient components to correct the velocity fields. We have also extended the AMR methodology to compressible viscous flows [17], with validation using carefully-designed shock tube experiments on unsteady shock reflections from wedges [14].

2.3 Software Framework

The technical goals of this require modeling of fluid flows in a broad range of conditions. Each regime requires specialized numerical methods particular to that regime. To manage the resulting code complexity we have developed each of the applications using a common computational framework. In particular, all of these codes are implemented in mixed-language paradigm of C++ and FORTRAN [13]. In this approach, programs are split into two pieces. The generation and manipulation of complex aggregate structures, such as the unions of rectangular grids that make up the AMR grid hierarchy, are defined and manipulated in C++, using the capability in that language for a user to define data types and operations on those data types. Floating-point intensive calculations, which are typically done on a single rectangular patch at a time, are implemented in FORTRAN. The C++ portions of these codes will be implemented using a data structure library specifically designed to support block-structured adaptive mesh calculations.

Our strategy for algorithm design is based on C++ class libraries to encapsulate the complex aggregate data structures used by our adaptive algorithms. This approach produces a natural paradigm for expressing parallelism. In particular, by developing parallel versions of the base class libraries that distribute the work, parallelism is essentially hidden from the physics packages that perform work on logically rectangular blocks of data. This approach has been successfully used for both shared memory parallelism using a work queue and on distributed memory machines using message passing. We are presently using the message passing system, MPI, for MIMD machines whose performance has been tuned for Cray T3E and IBM SP type machines.

In general, the main obstacles to parallelization are communications overhead and load balance. In gas dynamical simulations, the communications overhead is not significant because of the very large number of floating point operations required per cell per time-step. Load balance is a more challenging problem for AMR on distributed memory architectures because of the broad distribution of grid sizes in a given calculation. However, we have found it possible to achieve very good load balance as long as there are many more grids than processors. With an average of three grids per processor, we can achieve load balance efficiencies of greater than 95% [12, 16].

Our implementation employs the parallelization strategy that we term grid-based-parallelism. AMR algorithms represent a problem solution on a hierarchy of levels of refinement. On each individual level of refinement, the solution is represented by a large number of rectangular data arrays (grids), where any given grid represents the solution in a particular rectangular region of space. In 3D calculations, the number of grids on the finest level of refinement, where most of the work is done, can number several hundreds to several thousands. The number of computational cells in a given grid in 3D can range from 5000 to a quarter million. The grid is therefore an object

representing a very substantial amount of computational effort per time-step. We have found that this effectively hides any parallel overhead that may be required to manage the grid object.

3 Computational Example

A critical issue in the success of attacking underground chamber systems is the fate of chemical or biological agents within the chamber. Results from a DTRA field test, in which triethylene phosphate (TEP) was used a surrogate chemical agent, suggest a secondary detonation or deflagration occurring approximately three seconds after an explosion within a chamber system. One possible explanation of this secondary event and the associated pressure load is the ignition of TEP leaking from the ruptured tank. In order to evaluate this scenario we are performing computational studies to assess the dispersal of TEP within a chamber and its subsequent ignition and combustion.

As a first step in this process we are modeling the early blast environment within a chamber. The chamber is 32 meters long, 8 meters wide and 8 meters high. At one end of the chamber we place 13,000 gallons of TEP. At the other end of the chamber we place a 100 kg TNT charge. We simultaneously impact the TEP with a high velocity projectile which provides a simplified model for debris impact from the explosion. Field tests on projectiles impacting containers filled with liquid suggest that impact effectively destroys the container so that for this early prototype simulations we have ignored strength effects for the TEP container.

In Figure 2 we show time slices on cutting planes through the initial center of the TEP showing the mass of TEP per unit volume. For these figures the scale has been adjusted so that a density above .01 g/cc of TEP appears as red while the remainder of the palette shows TEP densities below this level. The time histories show that while the bulk of the TEP is falling to the floor and splashing along the walls within the first 1.43 seconds of evolution, there is, nevertheless, a residual low concentration haze of TEP that has been entrained into the surrounding air as a result of the impact and interaction with the blast waves from the explosion. It is this remnant haze that is mixed with the surrounding air that provides the potential fuel for the secondary deflagration/detonation that was observed during a field test. To quantify the available fuel we subdivide the chamber into 8 equally spaced vertical bins (one meter thick) and compute the distribution of TEP in these bins as function of time. These mass distributions are presented in Figure 3. The figure clearly shows that the bulk of the mass drops to the floor so that within the first .75 seconds most of the TEP mass is located in the lower two meters of the chamber. However, even at 1.43 seconds a substantial quantity of TEP remains in the upper layers. At this late time each of the upper layers contains over 600 kilograms of TEP, providing adequate fuel for secondary combustion.

This simulation was performed using 2 AMR levels of refinement, with a refinement ratio of 4 for each level. The base grid resolution of $32 \times 8 \times 8$ thus has an effective fine grid resolution of $512 \times 128 \times 128$: 8388608 computational zones distributed among approximately 250 grids. 32 nodes of an IBM SP2 system, pandeon, at ERDC were used, each node consisting of a single 160 MHz Power2 processor with 512 MB of memory. Approximately 5 minutes per time step were required.

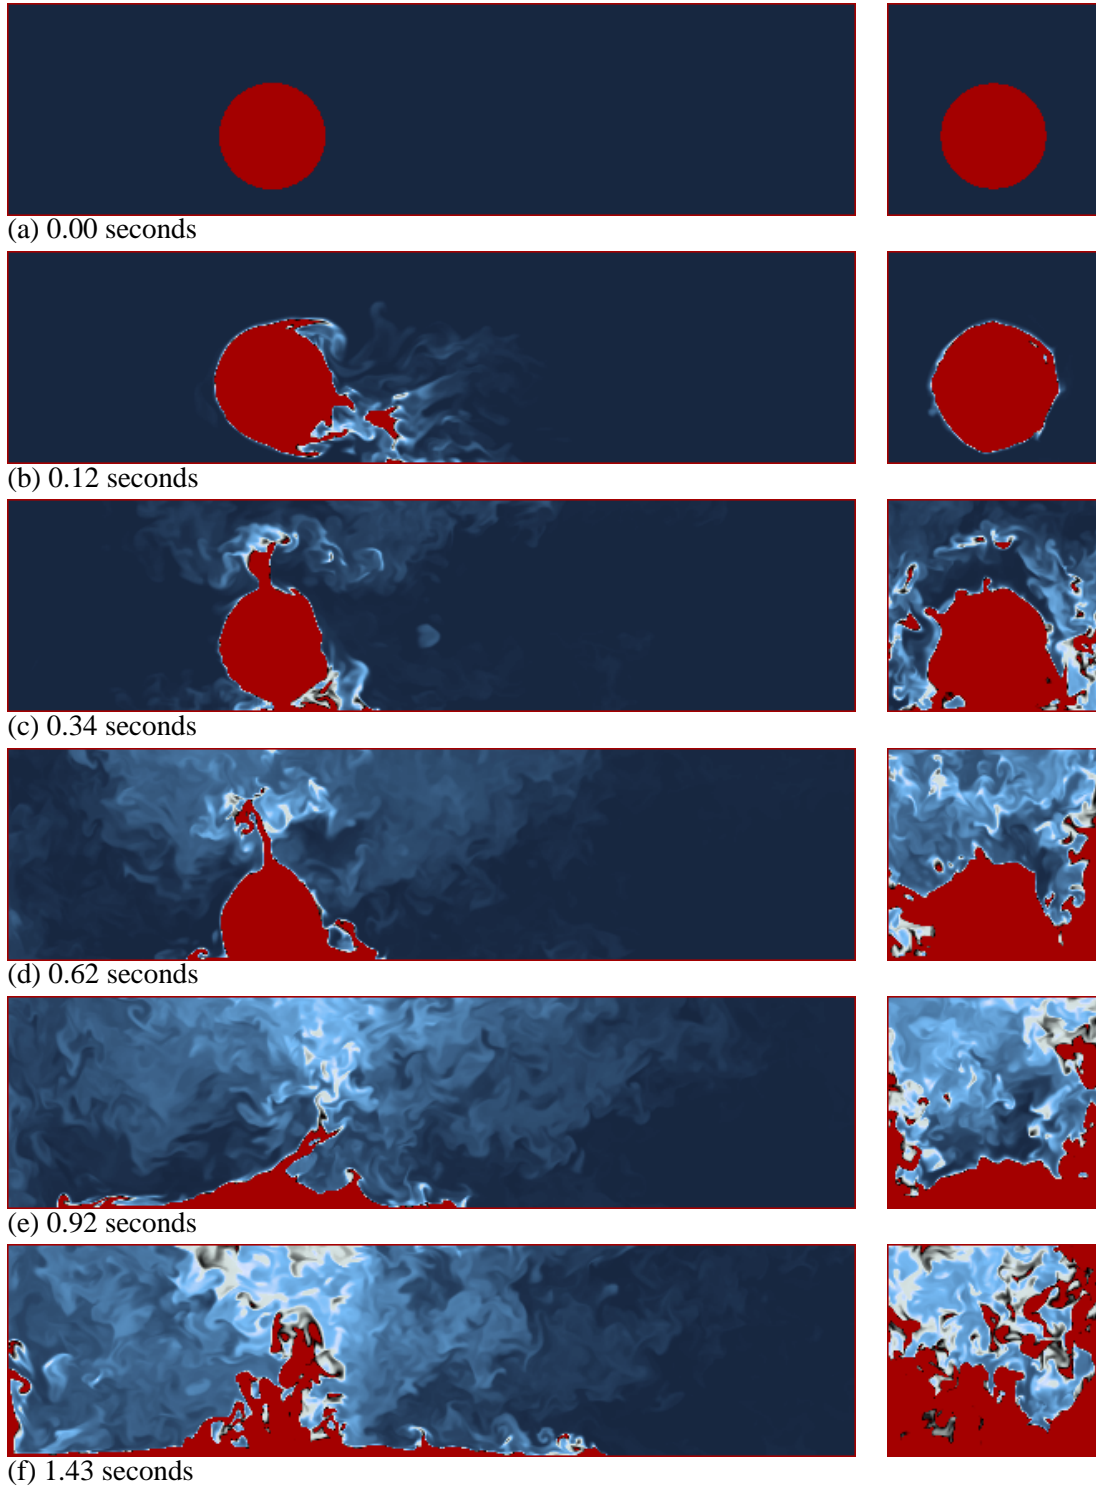


Figure 2: Evolution of mass of TEP per unit volume in vertical slices on planes through initial center of mass of TEP.

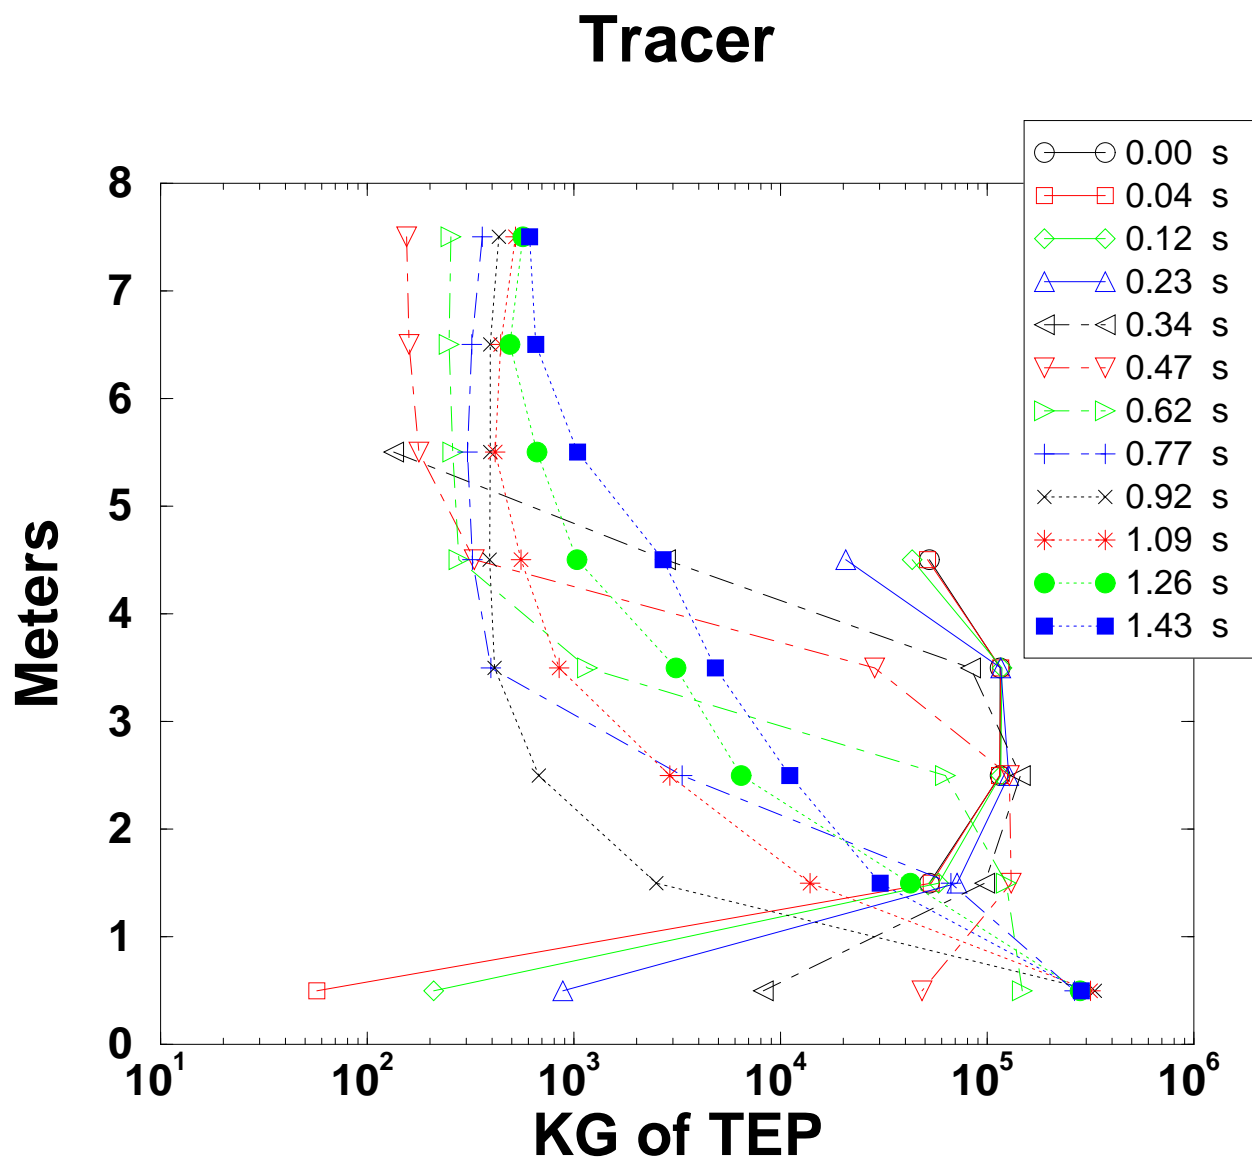


Figure 3: Vertical TEP mass distribution with increasing time.

4 Future Work

The computations performed to date provide a compelling qualitative picture to explain the observations obtained during field experiments. To further refine this model we plan to perform more detailed studies that incorporate various degrees of venting which will modify the pressure environment within the chamber and potentially exhaust TEP into the surrounding atmosphere. We also plan to incorporate a turbulent combustion model to quantitatively estimate the pressure rise that would be associated with deflagration of the TEP haze within the chamber. A critical component of this portion of the study will be developing suitable models for ignition of the mixture. Finally, we plan to investigate higher fidelity models for the thermochemical properties of TEP in order to provide a more realistic evaluation of this type of scenario.

References

- [1] A. S. Almgren, J. B. Bell, P. Colella, and L. H. Howell. Adaptive projection method for the incompressible Euler equations. In *Proceedings of the Eleventh AIAA Computational Fluid Dynamics Conference*, pages 530–539. AIAA, 1993.
- [2] A. S. Almgren, J. B. Bell, P. Colella, and L. H. Howell. An adaptive projection method for the incompressible Navier-Stokes equations. In *Proceedings of the IMACS 14th World Conference, Atlanta, GA, July 11–15, 1994*.
- [3] A. S. Almgren, J. B. Bell, P. Colella, L. H. Howell, and M. Welcome. A high-resolution adaptive projection method for regional atmospheric modeling. In *Proceedings of the U.S. EPA NGEMCOM Conference, August, 1995*.
- [4] J. Bell, M. Berger, J. Saltzman, and M. Welcome. A three dimensional adaptive mesh refinement for hyperbolic conservation laws. *SIAM J. Sci. Statist. Comput.*, 15(1):127–138, 1994.
- [5] J. B. Bell, P. Colella, and H. M. Glaz. A second-order projection method for the incompressible Navier-Stokes equations. *J. Comp. Phys.*, 85(2):257–283, 1989.
- [6] J. B. Bell, P. Colella, and J. A. Trangenstein. Higher order Godunov methods for general systems of hyperbolic conservation laws. *J. Comp. Phys.*, 92(2):362–397, 1989.
- [7] M. J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *J. Comp. Phys.*, 82(1):64–84, 1989.
- [8] M. J. Berger and J. Olinger. Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comp. Phys.*, 53:484–512, 1984.
- [9] A. J. Chorin. Numerical solution of the Navier-Stokes equations. *Mathematics of Computation*, 22:745–762, 1968.
- [10] P. Colella. Multidimensional upwind methods for hyperbolic conservation laws. *J. Comp. Phys.*, 87:171–200, 1990.

- [11] P. Colella and P. R. Woodward. The piecewise parabolic method (PPM) for gas-dynamical simulations. *J. Comp. Phys.*, 54:174–201, 1984.
- [12] W. Y. Crutchfield. Load balancing irregular algorithms. Technical Report UCRL-JC-107679, Applied Mathematics Group, Computing & Mathematics Research Division, Lawrence Livermore National Laboratory, July 1991.
- [13] W. Y. Crutchfield and M. L. Welcome. Object-oriented implementation of adaptive mesh refinement algorithms. *Journal of Scientific Programming*, 2:145–156, 1993.
- [14] L. F. Henderson, W. Y. Crutchfield, and R. J. Virgona. The effects of heat conductivity and viscosity of argon on shocks waves diffracting over rigid ramps. *J. Fluid Mech.*, 331:1–36, 1997.
- [15] M. L. Minion. On the stability of Godunov-projection methods for incompressible flow. *J. Comp. Phys.*, 123:435–449, 1996.
- [16] Charles A. Rendleman, Vincent E. Beckner, Mike Lijewski, William Crutchfield, and John B. Bell. Parallelization of structured, hierarchical adaptive mesh refinement algorithms. *Computing and Visualization in Science*, 3(4), 2001. accepted.
- [17] E. Steinthorsson, D. Modiano, W. Crutchfield, J. Bell, and P. Colella. An adaptive semi-implicit scheme for simulations of unsteady viscous compressible flows. In *Proceedings of the 12th AIAA Computational Fluid Dynamics Conference, San Diego, CA. June, 1995*. AIAA-95-1727-CP.
- [18] Bram van Leer. Towards the ultimate conservative difference scheme V: a second-order sequel to Godunov’s methods. *J. Comp. Phys.*, 32:101–136, 1979.